

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 9/48	A2	(11) International Publication Number: WO 00/67124 (43) International Publication Date: 9 November 2000 (09.11.00)
(21) International Application Number: PCT/US00/08157 (22) International Filing Date: 28 March 2000 (28.03.00) (30) Priority Data: 09/302,592 30 April 1999 (30.04.99) US (71) Applicant: HONEYWELL INC. [US/US]; 101 Columbia Road, Morristown, NJ 07692 (US). (72) Inventors: KENCHAMMANA-HOSEKOTE, Deepak, R.; 1180 Lochinvar Avenue #44, Sunnyvale, CA 94080 (US). VAIDYANATHAN, Nemmara, R.; 2790 Kelvin Avenue #1134, Irvine, CA 92614 (US). ALLALAGHATTA, Pavan; 545 Sandhurst Drive - #316, Roseville, MN 55113 (US). (74) Agent: MACKINNON, Ian, D.; Honeywell International Inc., Patent Services Group, AB-2B, P.O. Box 2245, Morristown, NJ 07962 (US).		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>Without international search report and to be republished upon receipt of that report.</i>
(54) Title: NETWORK SCHEDULER FOR REALTIME APPLICATIONS		
(57) Abstract		
<p>A network scheduler for real time applications is described. The network scheduler supports scheduling of network packets in a multi-threaded operating system to insure correct scheduling order and to preclude kernel preemption of active threads. The network scheduler resides between a kernel (and device drivers) and multiple sessions at the user-level. The network scheduler uses input/output control primitives such as HOLD and RELEASE to control servicing of the queues of packets at the kernel or network device driver. The control primitives serve as a middleware interface to the underlying kernel or network device driver. The network scheduler functions as a proxy scheduler to manage the queues of packets at the kernel or network device driver.</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

NETWORK SCHEDULER FOR REALTIME APPLICATIONS**FIELD OF THE INVENTION**

The present invention relates generally to real time applications and more specifically to a network scheduler for such applications.

COPYRIGHT NOTICE/PERMISSION

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever. The following notice applies to the software and data as described below and in the drawings hereto: Copyright 1998, 1999 Honeywell, Inc., All Rights Reserved.

STATEMENT OF GOVERNMENT INTEREST

This invention was made with Government support under contract number F19628-94-C-0044 awarded by Advanced Research Projects Agency (ARPA)/CSTO. The Government has certain rights in the invention.

BACKGROUND OF THE INVENTION

Real time systems comprise programs that must execute relative to priorities or within strict constraints on response time. Some examples of real-time systems include industrial plant controllers, medical monitors, communications networks, and aircraft or spacecraft flight control systems. One difficulty in designing real-time systems is in the strict requirement to meet processing deadlines and not in the actual speed of processing. A real-time system need not be complex or high performance. Checking radioactivity levels in a nuclear power plant or checking a pulse rate on a medical patient may require only simple processing. However, a failure to meet a processing deadline, such as decreasing plant power or sounding a medical alarm, can be literally fatal.

One example of a real time system is a multimedia system. Multimedia is increasingly recognized as an enabling technology with diverse applications from

interactive training to medical diagnosis to control applications, for example. In the context of control applications, the term multimedia implies a much broader scope than the traditional usage of the term that refers to audio, video and text. In the context of control applications, multimedia encompasses both the traditional multimedia data as well as image data and spectral data. Image data refers to still images, full-motion video, and graphics. Image data is obtained from sources such as video, X-ray, and scanning electron microscopy to name a few. Spectral data refers to sensor data. Spectral data is obtained from sources such as infrared, color, acoustic emission, ultrasound, microwave, laser, and optical sensors and devices, for example.

A multimedia-based control system imposes stringent requirements on a network. A key network requirement for multimedia-based control systems is that the underlying network must support real time tasks. Performance criteria for a network in such a control system are very different than those of a typical information network. Some basic performance criteria for a control network are (1) the guarantee of critical traffic, (2) the avoidance of priority inversion, and (3) the guarantee of deadlines. Priority inversion occurs when a lower priority task is allocated the resources before a higher priority task.

When scheduling real time tasks in a network, a priority or a deadline is associated with the task. For many real time tasks, failure to meet the deadline or failure to avoid priority inversion is considered a system fault. The priority or the deadline associated with a task is commonly specified globally at the user or application level. However, a kernel-level scheduling mechanism does not allow users or applications to exercise control over scheduling policies. As a result, user or application-level scheduling policies fail when kernel-level scheduling cannot enforce the same policy.

For these and other reasons, there is a need for a mechanism to guarantee correct scheduling order of real time tasks according to the user or application specified priorities or deadlines.

SUMMARY OF THE INVENTION

A network scheduler for real time applications is described. According to one aspect of the invention, a network scheduler acts as a proxy scheduler across different scheduling domains for the purpose of achieving correct deadline and priority based scheduling of network packets. An application specifies real time requirements in terms of deadlines or periodicities. An appropriate scheduling algorithm is applied based on the application task periodicities and deadlines. Simple primitives exercise control over operating system threads over scheduling domains to control their suspension and execution. The scheduling mechanism functions as a middleware system.

Another aspect of the invention is a scheduling method for an application distributed across a plurality of processing nodes. The scheduling method comprises a proxy scheduling mechanism at every node between the boundaries of a plurality of scheduling domains within a processing node. The scheduling method achieves overall end to end deadline and priority based scheduling for distributed real time applications running over a plurality of processing nodes with multiple scheduling domains.

The present invention is applicable to a variety of operating system platforms and middleware solutions for real time systems. The present invention allows user-level defined attributes, such as priorities and deadlines, and desirably overcomes constraints resulting from vendor specific operating system mechanisms that do not preserve such user specified attributes of real time behavior. The primitives of the invention allow various scheduling policies such as priority based or deadline based scheduling policies.

Still other and further embodiments, aspects and advantages of the invention will become apparent by reference to the drawings and by reading the following detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram illustrating a system level overview of an example embodiment of the invention.

FIG. 2 is a flowchart of an example embodiment of a method of performing priority based scheduling.

FIG. 3 is a flowchart of an example embodiment of a method of performing deadline based scheduling.

FIG. 4 is a diagram of a computerized system in conjunction with which example embodiments of the invention may be implemented.

FIG. 5 is a diagram of an example embodiment of a computer system incorporating a network scheduler block

5 FIG. 6 is a more detailed diagram of an example embodiment of a computer system incorporating a network scheduler block.

DESCRIPTION OF THE EMBODIMENTS

10 In the following detailed description of the embodiments, reference is made to the accompanying drawings that form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention, and it is to be understood that other embodiments may be utilized and that structural, logical and electrical changes may be made without departing from the scope of the present inventions. The following detailed description is, therefore, not
15 to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

Introduction. A network scheduler according to the present invention is implemented with a multithreaded operating system having different priority domains for user-level threads and kernel-level threads. A kernel-level scheduling mechanism
20 generally does not allow users to exercise control over scheduling policies. Thus, user level scheduling policies may fail if the kernel-level scheduling cannot enforce the same policy. The network scheduler solves these and other scheduling problems in a multithreaded operating system having different methods of scheduling threads within a user-level domain and a kernel-level domain. For example, in some multithreaded
25 operating systems a priority based scheduling policy or a deadline based scheduling policy is used at the user level and a round-robin based scheduling policy is used at the kernel level. In such an example, the network scheduler desirably prevents the kernel-level scheduling order from preempting the user-level scheduling order. As a result, the
30 network scheduler, as further described below, guarantees correct scheduling order according to the user specified priorities or deadlines.

The detailed description of this invention is divided into three sections. The first

section provides a system level overview of an example embodiment of the present invention. The second section describes methods of scheduling network communication packets according to an example embodiment of the invention. The third section describes a computing system in conjunction with which embodiments of the present invention are implemented.

System Overview. A system level overview of an example embodiment of the invention is described by reference to Figure 1. The network scheduling system 10, as shown in Figure 1, comprises one or more user-level sessions 12, a network scheduler 14, a kernel and device drivers 16 and communications hardware 18 including one or more physical communication channels.

The user-level sessions 12 have one or more user-level threads. Each thread is a different stream of control that can execute its instructions independently, allowing a multithreaded process to perform numerous tasks concurrently. Each one of the user-level threads is assigned a scheduling constraint based on the requirements of the real time application with which the user-level thread is associated. In one embodiment, the scheduling constraint is a priority. The priority is assigned using any common priority based scheduling policies. Two well known examples of priority based scheduling policies are Rate Monotonic Analysis and Deadline Monotonic Analysis. In an alternate embodiment, the scheduling constraint is a deadline. A common example of a deadline based scheduling policy is an Earliest Deadline First policy.

The network scheduler 14 schedules the use of a shared resource such as a physical communication channel. The network scheduler 14 receives a scheduling constraint for a user-level thread and manages the scheduling of the corresponding kernel-level thread based on the scheduling constraint of the user-level thread. The network scheduler ensures that kernel or device drivers 16 schedule the network packets in a manner that meets the scheduling constraint assigned by the user-level session 12. In one embodiment, the network scheduler uses instructions (referred to herein as primitives) to perform the following functions: hold a queue of packets, release a queue of packets, check the population of a queue, check a deadline of a packet at the head of a queue, and check the priority of a packet at the head of a queue. An example method of scheduling based on priorities is further described with respect to FIG. 2. An example

of scheduling based on deadlines is further described with respect to FIG. 3.

The kernel 16 (or any comparable center of an operating system) generally releases network communication packets according to a predetermined scheme. One example of a common scheme for kernel level scheduling is a round-robin scheme in which available communication packets are released in a cyclic manner. However, such a kernel-level scheduling policy may not enforce the scheduling constraints specified at the user-level session 12. The network scheduler 14 of the present invention solves this problem using the primitives described above to enforce the user specified scheduling constraints at the kernel and device driver level.

The device driver 16 maintains a queue of packets waiting to access the physical communication channel. The network scheduler 14 ensures that these packets are released from the queue according to the user specified scheduling constraints. In one embodiment, the network scheduler 14 is used to schedule the use of one or more physical communication channels between network components in a circuit switched network.

The process of determining which thread to allow to send packets over a network channel is referred to herein as scheduling. Scheduling for real time applications, such as multimedia applications, must be done within defined constraints or else the system may fail. In previous systems, the operating system controlled the scheduling of packets on the physical communication channel, but in a system according to the present invention the network scheduler is implemented in middleware and the network scheduler directly controls scheduling of packets using methods which are further described below. The example embodiment described below is directed to the use of the network scheduler to schedule network communication packets. However, one of skill in the art will recognize that the network scheduler can be adapted to schedule the use of other shared hardware or software resources.

Methods of the Invention. Referring now to FIGS. 2 and 3 a particular method performed by a processor in an exemplary embodiment is described by reference to a series of flowcharts. The methods to be performed constitute computer programs made up of computer-executable instructions. The methods shown in FIGS. 2 and 3 are implemented in a machine readable medium comprising machine readable instructions

for causing a computer to perform the method. Such machine readable medium may include software modules and computer programs. The type of computer programming languages used to write the code may vary from procedural code type languages to object oriented languages. The files or objects need not have a one to one
5 correspondence to the modules or method acts described depending on the desires of the programmer. Further, the method and apparatus may comprise combinations of software, hardware and firmware as is well known to those skilled in the art. Describing the methods by reference to flowcharts enables one skilled in the art to develop such programs including such instructions to carry out the methods on suitable computerized
10 systems.

FIG. 2 is a block diagram of an example embodiment of a method of performing priority based scheduling. In one embodiment, the network scheduler 14 of FIG. 1 acts as a proxy scheduler for the kernel by performing the method shown in FIG. 2. Using the method shown in FIG. 2, the network scheduler 14 uses a HOLD, a RELEASE, a
15 CHECK PRIORITY and a CHECK POPULATION primitive or the like to control the behavior of the kernel-level scheduler and to ensure that the user-level scheduling requirements are met.

In the embodiment of a priority based scheduling method 20 shown in FIG. 2, a
20 network scheduler checks for the arrival of a new packet for transmission on a network connection (block 22). Upon the arrival of a new packet, the network scheduler issues a HOLD primitive to hold all of the network connection queues (block 23). The HOLD primitive holds the connection queues while the priority of the packet is checked. The network scheduler checks the priority assigned to the arrived packet at the user level
25 (block 24). If no priority is assigned to the arrived packet, the packet is assigned to a general queue. In addition to checking the priority of the arrived packet, the priority of the arrived packet is compared to a priority level of a highest priority queue (block 26.) If the priority of the arrived packet is less than the priority of the queue (block 28), then the network scheduler issues a CHECK POPULATION primitive to determine if the
30 queue has packets pending (block 30.) If the queue is populated with pending packets (block 32), a RELEASE primitive is issued by the network scheduler to release the packet(s) in the queue (block 34.) The network scheduler releases the queue at block 34

because the queue has a user specified priority that is a higher priority than the user specified priority of the arrived packet. After the queue is released, the priority level of the arrived packet is compared with the priority level of the next highest level priority queue (block 26.) The actions in blocks 26, 28, 30, 32, and 34 are repeated until the
5 priority level of the packet is not less than the priority level of the queue.

If the priority level of the arrived packet is not less than the priority level of the queue, then the priority level of the arrived packet is equal to the priority level of the queue (block 36.) The network scheduler issues a HOLD primitive for the queue that has a priority level equal to the arrived packets priority level (block 38.) The arrived
10 packet is assigned to the queue with the same priority level (block 40) and then the network scheduler checks for the arrival of another new packet (block 22.)

If no new packets have arrived (block 22), then the queues are checked beginning with the highest priority level queue (block 42). The network scheduler issues a RELEASE primitive for the queue (block 44). After releasing the queue, the network
15 scheduler checks for the arrival of a new packet (block 46.) If a new packet has not arrived, the queue with the next highest level priority is checked (block 42) and the actions in blocks 42, 44, and 46 are repeated until a new packet arrives.

An example embodiment of a method of performing priority based scheduling has been described by reference to FIG. 2. In an alternate embodiment, the network
20 scheduler 14 of FIG. 1 acts as a proxy scheduler for the kernel by performing the method shown in FIG. 3.

FIG. 3 is a block diagram of an example embodiment of a method of performing deadline based scheduling. Using the method shown in FIG. 3, the network scheduler
25 14 uses the HOLD, RELEASE, CHECK POPULATION and CHECK DEADLINE primitives to control the behavior of the kernel-level scheduler and to ensure that the deadline based user-level scheduling requirements are met.

The deadline based scheduling method 60 shown in FIG. 3, checks for the arrival of a new packet for transmission on a network connection (block 62). Upon the arrival
30 of a new packet, the network scheduler issues a HOLD primitive to hold all of the network connection queues (block 64). The HOLD primitive holds the connection queues while the deadline of the arrived packet is checked. The network scheduler

checks the deadline assigned at the user level to the arrived packet (block 66).

As well as checking the deadline of the arrived packet, the deadline of the arrived packet is compared to the deadline of the packet at the head of the queue with the earliest deadline (block 68). If the deadline of the queue is earlier than the deadline of the arrived packet (block 72), the queue is released using the RELEASE primitive (block 74). After releasing the queue, the deadline of the arrived packet is compared to the deadline of the queue with the next earliest deadline (block 68). The actions in block 68, 72, and 74 are repeated until the deadline of the arrived packet is earlier than the deadline of the next queue. Since the deadline of the queue is later than the deadline of the arrived packet (block 72), the arrived packet is inserted to an unassigned queue (block 76) and a new packet is checked for (block 62).

If no new packets have arrived (block 62), the next held queue is checked in order deadlines beginning with the earliest deadline (block 78). The network scheduler issues a RELEASE primitive for the queue having the earliest deadline (block 80) and the network scheduler again checks for the arrival of a new packet (block 62).

An example embodiment of a method of performing deadline based scheduling has been described by reference to FIG. 3. One skilled in the art will recognize that a data structure maintains an ordered list of the queues (maintained in deadline order) for the example embodiment listed in FIG. 3. When a new packet is assigned to a queue or when a queue is released, the data structure is updated according to any means of checking, updating or re-ordering such a data structure.

Hardware and Operating Environment. FIG. 4 is a diagram of a computerized system in conjunction with which example embodiments of the invention may be implemented. Computer 90 includes a processor 92, random-access memory (RAM) 94, read-only memory (ROM) 96, and one or more storage devices 98, such as a hard disk drive, a floppy disk drive (into which a floppy disk can be inserted), an optical disk drive, a tape cartridge drive and the like. RAM 94 and ROM 96 are collectively referred to as the memory of computer. The memory, hard drives, floppy disks, and the like, are types of computer-readable media. The computer-readable media store instructions for execution by the processor 92. The instructions are provided to the processor 92 over a bus. The invention is not particularly limited to any type of computer 90. The

construction and operation of such computers are well known within the art.

In one embodiment, the network scheduler is incorporated into software executing on a computer, such as the computer 90. As described above, the software program guarantees correct scheduling order according to the user specified priorities or deadlines. The software program is developed using any high level language. In an

example embodiment, the software is developed using the C++ programming language. In one embodiment, the software executes under an operating system such as Solaris 2.5 or the like and executes on an industry-standard workstation such as a Sun SPARCstation 20 or the like connected on a Fore Systems ATM (Asynchronous Transfer Mode) LAN.

Those skilled in the art will recognize, however, that other comparable hardware and operating systems may be employed without diverting from the scope of the present invention. For example, the network scheduler may be employed with any operating system having different scheduling domains at the user level and at the kernel level as well as different methods of priority implementations and servicing of threads within the user domain and the kernel domain.

An alternate embodiment of the network scheduler of the present invention is implemented in a distributed block based programming environment such as described in a co-assigned US patent application filed on December 31, 1998, entitled System for Constructing and Distributing Block-Based Program Fragments [Attorney docket number H16-16875(256.015us1)]. The section of that co-assigned application entitled Distributed Block-Based Application Programming Model and the section entitled System Overview are herein incorporated by reference. In one embodiment, the network scheduler of the present invention operates as a system-level block in such a distributed block-based application programming model.

FIG. 5 is a diagram of an example embodiment of a computer system incorporating a network scheduler block according to a block based programming model. The computer system of FIG. 5 comprises two computers 110, 112, connected to a network 114. One or more distributed block based application programs 116 run on the computer. A network scheduler block 118 uses primitives, such as the primitives described earlier, to exercise control over operating system threads responsible for

scheduling the release of network packets. As shown in FIG. 5, each computer 110, 112 (or processing node) has a network scheduler block acting as a proxy scheduling mechanism between the boundaries of different scheduling domains. The computer system of FIG. 5 executes a distributed multimedia control application for displaying
5 images recorded by a video camera on a display.

FIG. 6 is a more detailed diagram of an example embodiment of a computer system incorporating a network scheduler block 118. In FIG. 6, the network scheduler block 118 acts as a proxy scheduler for packets on an ATM network 130. As shown in FIG. 6, one or more distributed block based application programs 102 execute on a
10 computer or processing node. Network communication packets associated with the application programs 102 are held in one or more queues 104. The network scheduler block 118 preserves scheduling attributes specified by the application programs 102. In the example embodiment shown in FIG. 6, the network scheduler block 118 schedules
15 packets for a switched virtual circuit 128 over the ATM network 130.

It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reviewing the above description. The scope of the invention should, therefore, be
20 determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

What is claimed is:

1. A method of scheduling the use of a resource of a computer system, the method
5 comprising:
 receiving a scheduling constraint for a user-level thread; and
 managing the scheduling of a corresponding kernel-level thread based on the
scheduling constraint of the user-level thread.
- 10 2. The method of claim 1, wherein the act of managing the scheduling of the kernel-
level thread further comprises issuing one or more primitives to control scheduling of
the kernel-level thread.
- 15 3. The method of claim 2, wherein the one or more primitives are selected from the
group consisting of: hold, release, check population of queue, check deadline of packet
that is at the head of the queue, and check priority of packet that is at the head of the
queue.
- 20 4. The method of claim 2, wherein the scheduling constraint is a priority assigned to
the user-level thread.
5. The method of claim 4, wherein the scheduling policy is a rate monotonic
analysis.
- 25 6. The method of claim 4, wherein the scheduling policy is a deadline monotonic
analysis.
7. The method of claim 2, wherein the scheduling constraint is a deadline assigned
to the user-level thread.
- 30 8. The method of claim 7, wherein the scheduling policy is an earliest deadline first
policy.

9. The method of claim 2, wherein the resource is a physical network communication channel.

5 10. A computerized system comprising:

an application program operating in a first scheduling domain, wherein the application program specifies real time scheduling requirements for a plurality of network packets;

an operating system operating in a second scheduling domain; and

10 a middleware scheduler for managing scheduling of the plurality of network packets in the second scheduling domain according to the scheduling constraints assigned in the first scheduling domain.

11. The computerized system of claim 10, wherein the real time scheduling
15 requirements are based on deadlines.

12. The computerized system of claim 10, wherein the real time scheduling requirements are based on priorities.

20 13. The computerized system of claim 10, wherein the application program is a block based application program.

14. The computerized system of claim 13, wherein the block based application program is distributed across a network.

25

15. A computer-readable medium having computer-executable instructions stored thereon for executing a method of scheduling network packets, the method comprising:

holding a plurality network connection queues;

checking a deadline of a first one of the queues;

30 comparing the deadline of the first one of the queues to a deadline of a new network packet;

releasing the first one of the queues if the deadline of the first one of the queues is

earlier than the deadline of the new network packet; and

assigning the new packet to a second one of the queues if the deadline of the first one of the queues is later than the deadline of the new network packet.

5 16. The computer-readable medium of claim 15 further comprising checking the first one of the queues to determine if the first one of the queues is populated.

17. A computer-readable medium having computer-executable instructions stored thereon for executing a method of scheduling network packets, the method comprising:

10 holding a plurality network connection queues;

checking a priority of a first one of the queues;

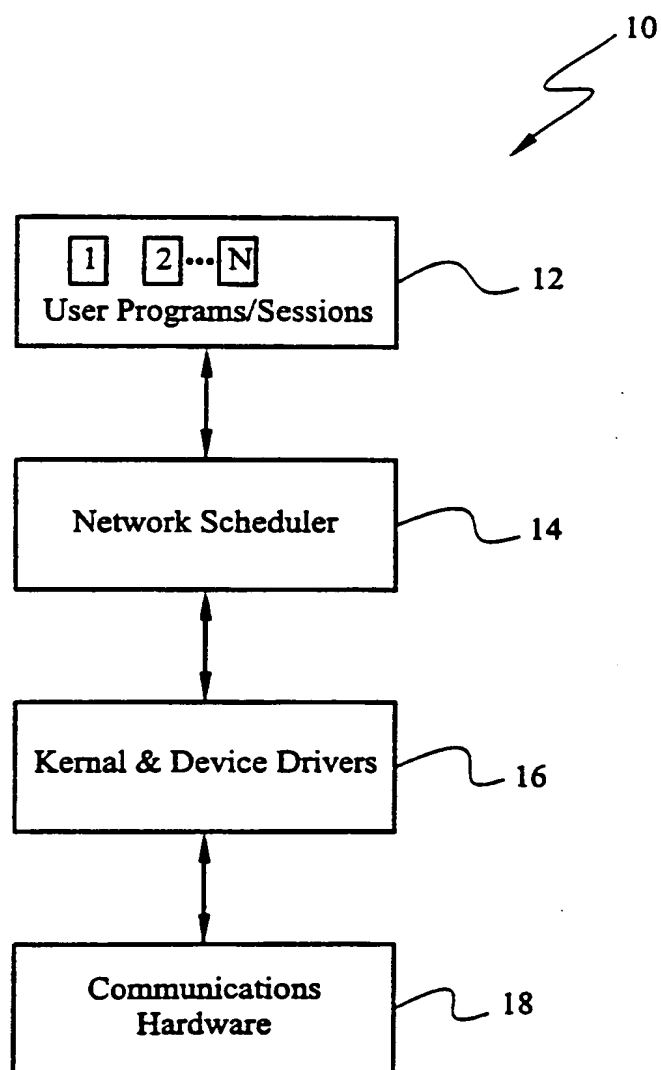
comparing the priority of the first one of the queues to a priority of a new network packet;

15 releasing the first one of the queues if the priority of the first one of the queues is greater than the priority of the new network packet; and

assigning the new packet to a second one of the queues if the priority of the first one of the queues is less than the deadline of the new network packet.

20 18. The computer-readable medium of claim 17 further comprising checking the first one of the queues to determine if the first one of the queues is populated.

1/6

*Fig. 1*

2/6

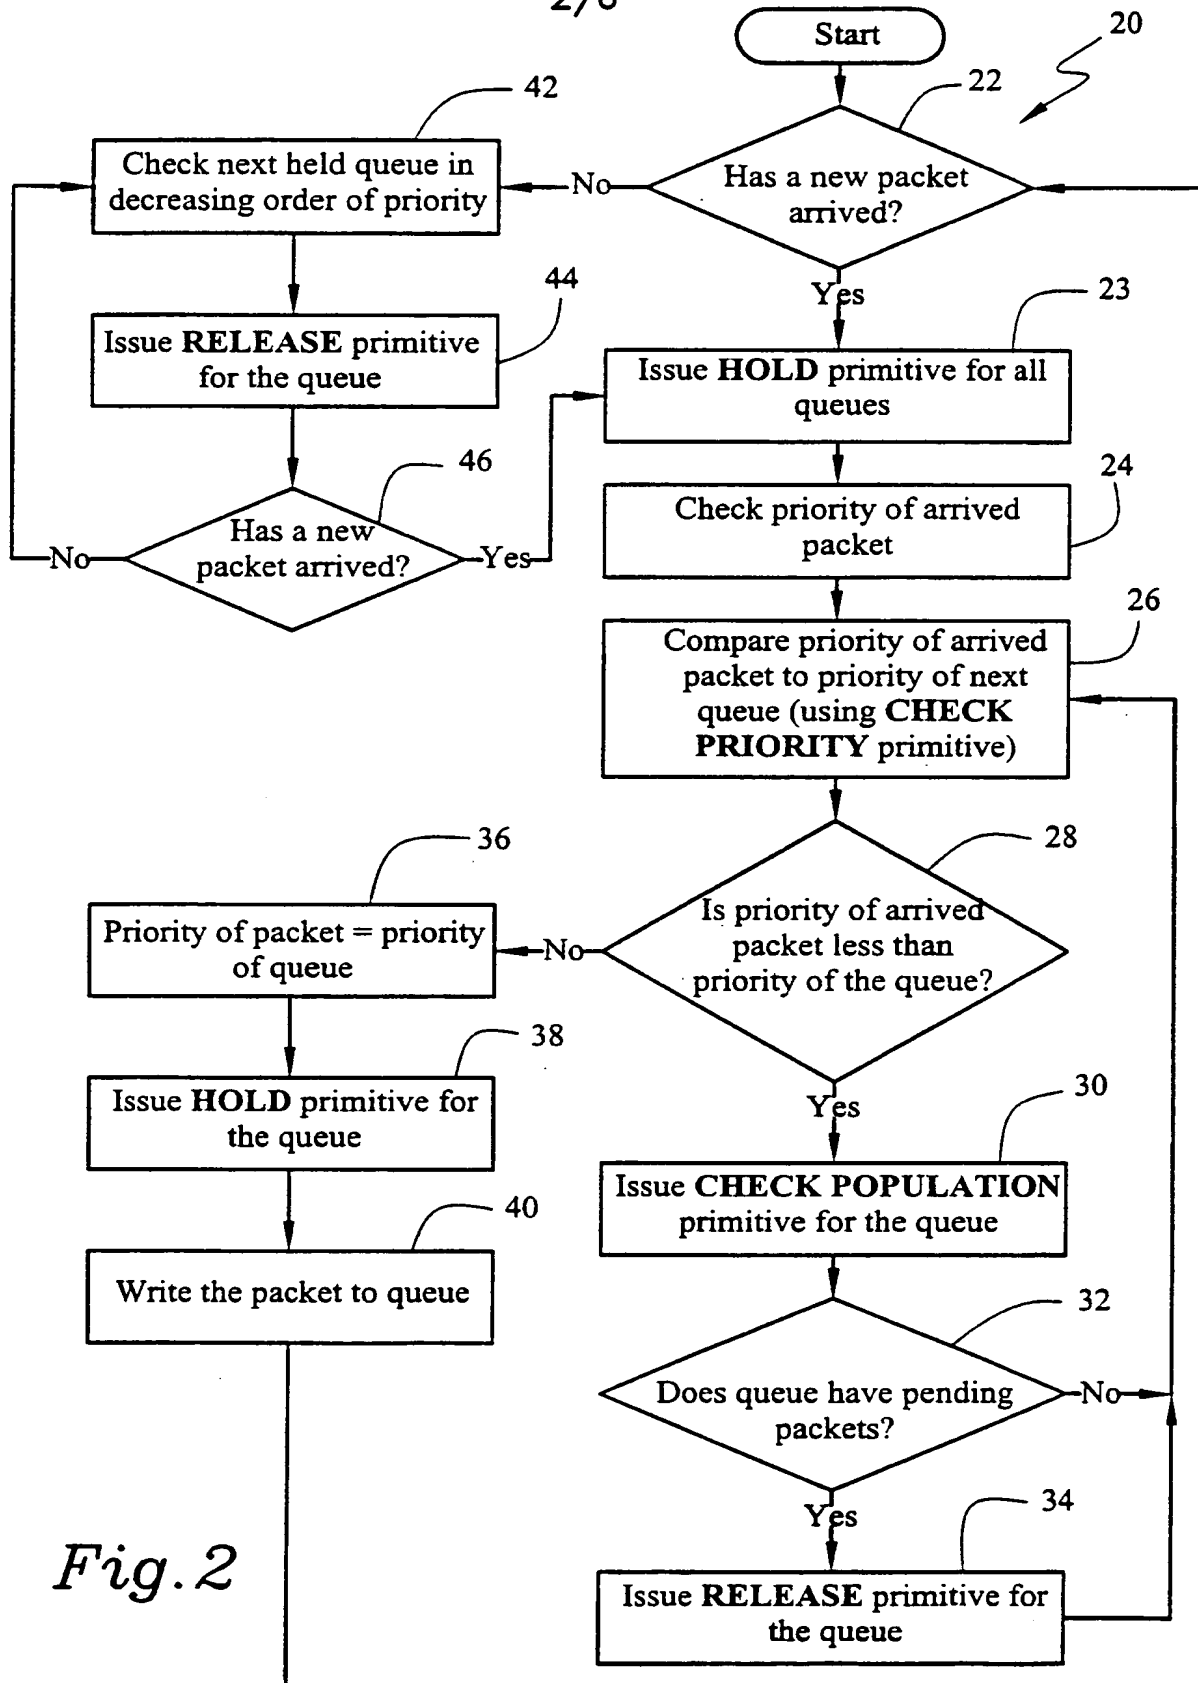
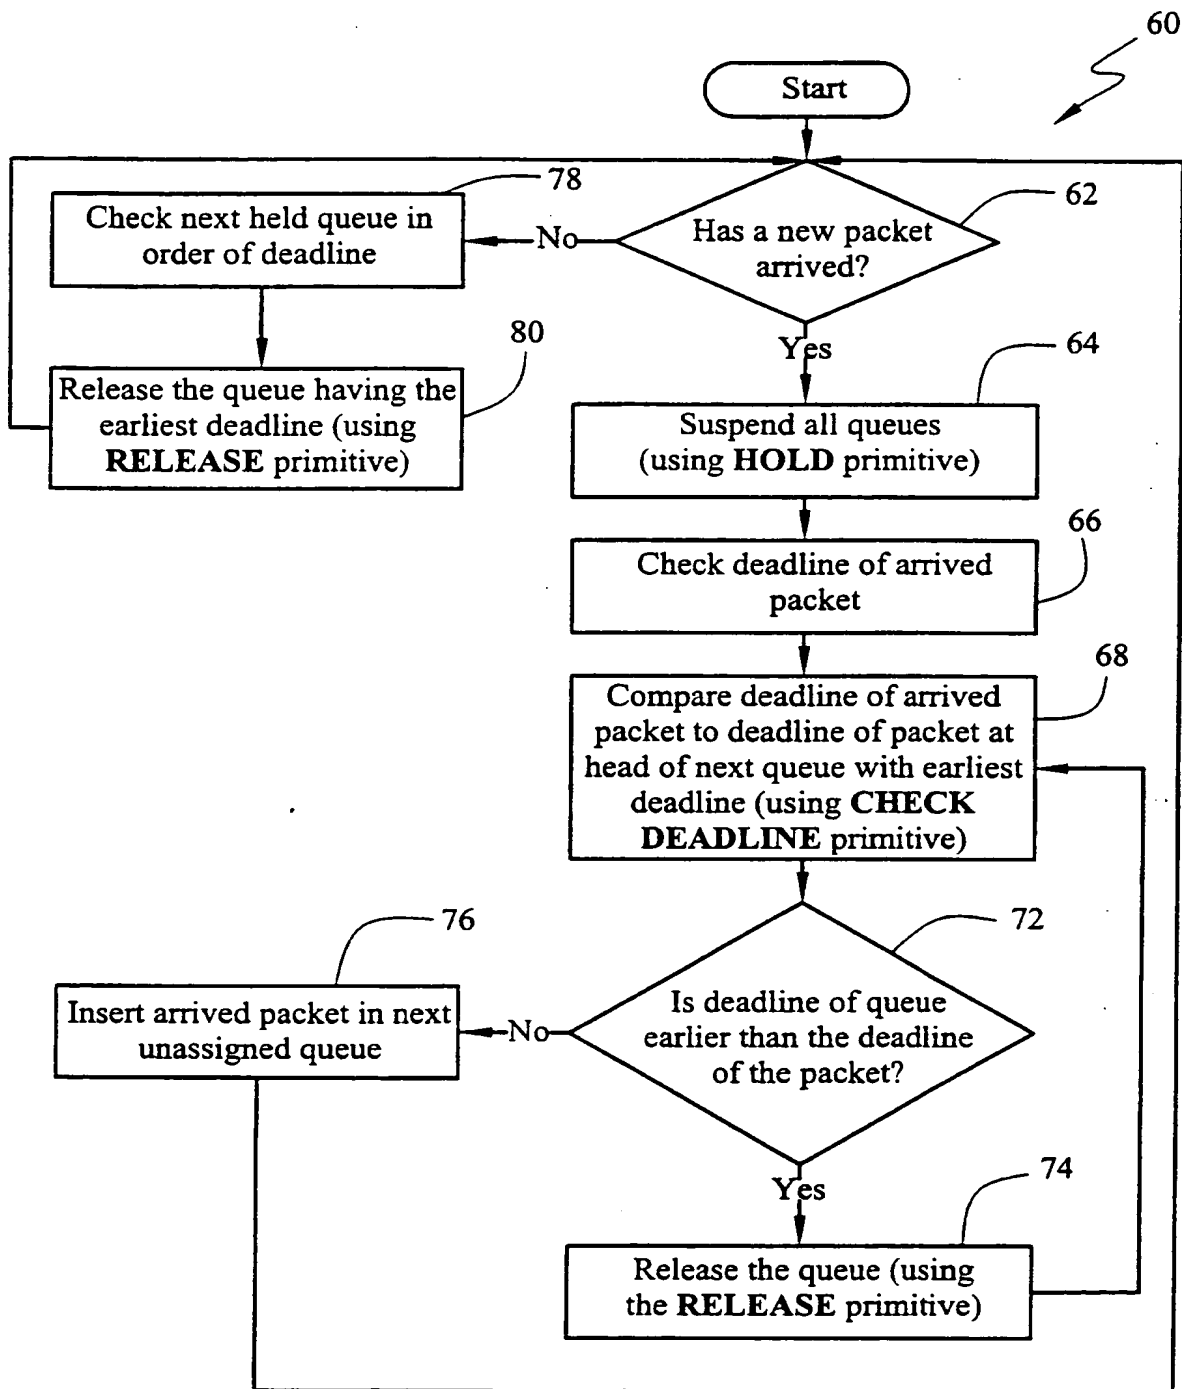


Fig. 2

3/6

*Fig. 3*

4/6





Fig. 4

5/6



6/6




Fig. 6

Docket # P2001,0349

Applic. # _____

Applicant: ALEXANDER BENEDIX ET AL.

Lerner and Greenberg, P.A.

Post Office Box 2480

Hollywood, FL 33022-2480

Tel: (954) 925-1100 Fax: (954) 925-1101